

L_1 Projections with Box Constraints

Mithun Das Gupta

Epson Research and Development, Inc.
San Jose, CA.

Sanjeev Kumar

Dept. Elec and Comp Engg.
University of California, San Diego

Jing Xiao

Epson Research and Development, Inc.
San Jose, CA.

Abstract

We study the L_1 minimization problem with additional box constraints. We motivate the problem with two different views of optimality considerations. We look into imposing such constraints in projected gradient techniques and propose a worst case linear time algorithm to perform such projections. We demonstrate the merits and effectiveness of our algorithms on synthetic as well as real experiments.

1 Introduction

In the domain of constrained optimization, it is well understood that L_2 norm penalty imposes smoothness constraint while the L_1 norm imposes sparsity [12]. Lately, sparse representations have been shown to be extremely efficient in encoding specific kinds of data, mainly, obeying power decay law in some transform space e.g. DCT etc. Donoho [6] provided sufficient conditions for obtaining an optimal L_1 -norm solution which is sparse. Recent work on compressed sensing [3, 16] further explores how L_1 constraints can be used for recovering a sparse signal sampled below the Nyquist rate.

L_1 regularized maximum likelihood can be cast as a constrained optimization problem. Although standard algorithms such as interior-point methods [15, 9] offer powerful theoretical guarantees (e.g., polynomial-time complexity, ignoring the cost of evaluating the function), these methods typically require at each iteration the solution of a large, highly ill-conditioned linear system; which are potentially very difficult and expensive.

This paper explores the behavior of L_1 constraint optimization under the presence of bound constraints. Traditional L_1 constraint assumes an infinite upper bound on the magnitude of the predicates. The idea of upper bounds can be explained very easily by a very simple example. Suppose we want to send a signal of length n which is a combined

signal originating from k different sources. Suppose at the receiver side we have k sets of receivers to decode the entire signal of length n . Based on the receiver set $i \in \{1, k\}$, the peak signal strength which the receivers can handle can be different. This kind of problem, is not handled by traditional L_1 projection. An illustration for the proposed problem, with 2 sets of receivers is shown in Fig. 1. Assuming that transmitting 0's or upper bounds cost just 1 bit, the effective sparsity, assuming 8 bits per element, is $(5*8+3*1)/64$ for L_1 , and $(2*8+6*1)/64$ for upper bounded L_1 . We motivate the inclusion of box constraints by look-




Figure 1: Left: L_1 projection and right: upper bounded L_1 projection for the same norm bound. Two colors represent different set of receivers with different upper bounds represented by black lines.

ing at the problem from two different settings.

Optimality Gap: Let us consider an unconstrained prob-

lem

$$\min_{\mathbf{x}} \|\mathbf{x} - \mathbf{v}\|_2^2 + \sum_i \lambda_i |x_i| + \zeta^T(\mathbf{1} - \mathbf{x}) + \gamma^T(\mathbf{x} - \mathbf{b})$$

where we have introduced the two bound constraints $[\mathbf{1}, \mathbf{b}]$ into the cost function. This can be slightly modified to a function of 2 variables such that

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{z}} \quad & \mathbf{z}^T \mathbf{z} + \sum_i \lambda_i |x_i| + \zeta^T(\mathbf{1} - \mathbf{x}) + \gamma^T(\mathbf{x} - \mathbf{b}), \\ \text{s.t.} \quad & \mathbf{z} = \mathbf{x} - \mathbf{v} \end{aligned}$$

The Lagrangian for this problem can now be written as

$$\begin{aligned} L(\mathbf{x}, \mathbf{z}, \beta) = & \mathbf{z}^T \mathbf{z} + \sum_i \lambda_i |x_i| + \zeta^T(\mathbf{1} - \mathbf{x}) \\ & + \gamma^T(\mathbf{x} - \mathbf{b}) + \beta^T(\mathbf{x} - \mathbf{v} - \mathbf{z}) \end{aligned}$$

The dual function is given by

$$\begin{aligned} \inf_{\mathbf{x}, \mathbf{z}} L(\mathbf{x}, \mathbf{z}, \beta) &= \inf_{\mathbf{z}} (\mathbf{z}^T \mathbf{z} - \beta^T \mathbf{z}) + \\ & \inf_{\mathbf{x}} (\sum_i \lambda_i |x_i| + (\beta - \zeta + \gamma)^T \mathbf{x}) \\ & - \beta^T \mathbf{v} + \zeta^T \mathbf{1} - \gamma^T \mathbf{b} \\ = & \begin{cases} -\frac{\beta^T \beta}{4} - \beta^T \mathbf{v} + \zeta^T \mathbf{1} - \gamma^T \mathbf{b} & \text{if } (\beta - \zeta + \gamma)_i \leq \lambda_i \\ -\infty & \text{otherwise} \end{cases} \end{aligned}$$

The Lagrange dual can now be written as

$$\begin{aligned} \max \quad & G(\beta) = -\frac{\beta^T \beta}{4} - \beta^T \mathbf{v} + \zeta^T \mathbf{1} - \gamma^T \mathbf{b} \\ \text{s.t.} \quad & (\beta - \zeta + \gamma)_i \leq \lambda_i \end{aligned}$$

A change of variables $\mu = (\beta - \zeta + \gamma)$ leads to

$$\begin{aligned} \max_{\mu_i \leq \lambda_i} G(\mu) &= -\frac{\mu^T \mu}{4} - \mu^T \mathbf{v} - \frac{\mu^T (\zeta - \gamma)}{2} \\ & - \frac{(\zeta - \gamma)^T (\zeta - \gamma)}{4} \\ & - (\zeta - \gamma)^T \mathbf{v} + \zeta^T \mathbf{1} - \gamma^T \mathbf{b} \end{aligned}$$

Now the duality gap for the bounded problem can be written as

$$\begin{aligned} \eta &= \|\mathbf{x} - \mathbf{v}\|_2^2 + \sum_i \lambda_i |x_i| + \\ & \zeta^T(\mathbf{1} - \mathbf{x}) + \gamma^T(\mathbf{x} - \mathbf{b}) - G(\mu) \\ &= \eta_{L1} - (\zeta - \gamma)^T(\mathbf{x} - \mathbf{v}) + \\ & \frac{(2\mu + (\zeta - \gamma))^T (\zeta - \gamma)}{4} \end{aligned}$$

where η_{L1} is the duality gap for the L_1 problem without the bound constraints. For x_i fixed at its upper bound b_i , $\gamma > \mathbf{0}$ and $\zeta = \mathbf{0}$. Under such a condition

$$\eta = \eta_{L1} - \gamma^T(\mathbf{v} - \mathbf{b}) - \frac{(2\mu - \gamma)^T \gamma}{4}$$

As long as $\gamma^T(\mathbf{v} - \mathbf{b}) + \frac{(2\mu - \gamma)^T \gamma}{4} \geq 0$ the duality gap is reduced as compared to simple L_1 minimization. Since $\gamma^T(\mathbf{v} - \mathbf{x})$ is always positive (shown later in Sec. 3), a sufficient condition for reduction in duality gap is $\mu \geq \frac{\gamma}{2}$, where μ is the dual feasible solution for the L_1 problem. The optimality of a particular solution is based on the duality gap and as such any decrement of the gap increases the optimality of the solution obtained. This shows that the optimality gap for the bound constrained L_1 problem can be made arbitrarily closer to zero compared to the similar unbounded L_1 problem.

Degrees of freedom for upper bounded problem: We study the degrees of freedom of the upper bounded L_1 projection problem in the framework of Stein's unbiased risk estimation (SURE) [14]. As shown by Zou et al. [17] the number of non-zero coefficients is an unbiased estimate for the degrees of freedom (DF) of the optimization scheme. The idea of upper bounds can be explained very easily by a very simple example. Suppose we want to send a signal of length n which is a combined signal originating from k different sources. Suppose at the receiver side we have k sets of receivers to decode the entire signal of length n . Based on the receiver set $i \in \{1, k\}$, the peak signal strength which the receivers can handle can be different. This kind of problem, is not handled by traditional L_1 projection. Assuming that the lower bound is zero (e.g. for electrical signals), transmitting upper bounds can cost just 1 bit. The effective sparsity, can be improved by sending 1 bit for all the elements fixed at their corresponding bounds, and sending 8 bit real numbers for all the remaining non-zero entities as compared to sending 8 bit reals for all non-zero entities. We conjecture that the degree of freedom for the bounded L_1 problem is bounded below that for the unbounded L_1 problem, and hence can provide increased sparsity in terms of the bounds.




Figure 2: SURE criteria for L_1 (blue) and upper bounded L_1 (red).

Let us again assume the simple estimation problem $y = \Phi x + \epsilon$. Now for estimating the SURE criterion, we estimate x and then generate the a new set of observation \hat{y} . The covariance between the terms y and \hat{y} is a scaled measure for the DF. Fig. 2 shows the covariance estimate for L_1 compared to upper bounded L_1 . Upper bounded L_1 is

uniformly bounded below the DF for only L_1 . This also emphasizes the conjecture for bounded constrained L_1 , the predicates which are fixed at their consecutive upper bound can be considered to fixed and as such do not contribute to the model complexity.

2 Separable Quadratic Problems

The problem of separable quadratic programming with linear bound constraints was first considered by Megiddo et al. [11]. They proposed linear time solution to the generic problem by Lagrangian relaxation based on the multidimensional search procedure of [10]. We introduce a novel linear time algorithm for gradient projection based norm minimization problem. Our starting point is an efficient method for projection onto the probabilistic simplex, with additional upper bound constraints. For infinite upper bound, this is the same method as proposed by numerous authors, namely Gafni et al. [8], Bertsekas [1], Crammer et al. [5], and more recently by ShalevSchwartz et al. [13] and Duchi et al. [7]. The basic intuition is that once the vector to be projected is ordered then the projection can be calculated exactly in linear time. Duchi et al. [7] proved the similarity of L_1 projection to the simplex projection, although the problem can be traced back to a special case of separable quadratic problem tackled by Megiddo et al. [11]. Although the projection step is linear, the sorting/ordering step is still $O(n \log n)$. We contest that once the upper bounds are introduced, the ordering does not remain as simple as the previous works. In this paper we propose a linear time algorithm which orders the difference of the bounds from the gradient vector to compute the projections on the norm constraint.

3 Upper bounded Simplex Projection

The most basic projection task we consider can be formally described as the following optimization problem, for $\mathbf{v} \in \mathbb{R}_+^n$

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \text{ s.t. } \{\mathbf{x} \in \Omega : \sum_{j=1}^n x_j \leq z, \mathbf{0} \preceq \mathbf{x} \preceq \mathbf{b}\} \quad (1)$$

where $\mathbf{0} \in \mathbb{R}^n$ is the vector of all zeros, $\mathbf{b} \in \mathbb{R}_+^n$ is the vector of upper bounds and $(\mathbf{p} \preceq \mathbf{q})$ denotes that $p_i \leq q_i, \forall i \in \{1, n\}$. Note that we enforce $\sum_{j=1}^n v_j \geq z$, because otherwise $\mathbf{x} = \mathbf{v}$ is the optimal solution. We assume that the feasibility of the constraints with respect to each other, and the existence of a solution is guaranteed by the set Ω being non-empty. Also note that if $z\mathbf{1} \preceq \mathbf{b}$ or $\mathbf{v} \preceq \mathbf{b}$, where $\mathbf{1}$ is the vector of all ones of size n , then the problem reduces to the projection onto the simplex problem of [7].

Claim 1. $x_i \leq v_i, \forall i \in \{1, 2, \dots, n\}$.

The norm of \mathbf{v} is larger than the constraint z as mentioned earlier. Assume there is an optimal projection \mathbf{x}^* , with its element $x_i^* > v_i$. There exists another solution $\hat{\mathbf{x}}^*$, such that $\hat{x}_i^* = v_i$ and all other elements same as \mathbf{x}^* , which is bounded and gives a lower value for the cost function, and hence is a better solution than \mathbf{x}^* , which is a contradiction.

Claim 2. $b_i^{\text{effective}} = \min(b_i, v_i), \forall i \in \{1, 2, \dots, n\}$.

The maximum value that x_i can reach is either the upper bound b_i or v_i (from Claim 1), hence the effective upper bound $b_i^{\text{effective}}$ is the minimum of b_i or v_i .

From this point onwards, we will assume $b_i = b_i^{\text{effective}} \forall i$, if not mentioned otherwise. The next lemma is the extension of Lemma 1 of Duchi et al. [7] for bounded projections.

Lemma 1. *If $v_i > v_j$ and in optimal solution $x_i = 0$ then $x_j = 0$, irrespective of ordering of b_i, b_j .*

Proof. We need to prove the above lemma for 2 cases.

[1] $b_i \geq b_j$. In this case, if another solution is constructed such that x_i and x_j are switched keeping all other indices same, then there is a strict decrease in optimal value, generating a contradiction.

[2] $b_i < b_j$. Again let us assume $x_j > 0$. Let us construct another optimal solution $\hat{\mathbf{x}}$, such that $\hat{x}_i = \Delta, \hat{x}_j = x_j - \Delta$, where $\Delta = \min(b_i, x_j)$, and keep all the other indices same. It can be easily observed that the norm as well as the upper bound constraint are satisfied for $\hat{\mathbf{x}}$. Now we can show that

$$\begin{aligned} & \text{New obj value} - \text{Old obj value} \\ &= (v_i - \hat{x}_i)^2 + (v_j - \hat{x}_j)^2 - (v_i - x_i)^2 - (v_j - x_j)^2 \\ &= (v_i - \Delta)^2 + (v_j - x_j + \Delta)^2 - v_i^2 - (v_j - x_j)^2 \\ &= \Delta^2 - 2v_i\Delta + \Delta^2 + 2v_j\Delta - 2x_j\Delta \\ &= 2\Delta(\Delta - x_j) + 2(v_j - v_i)\Delta \leq 2(v_j - v_i)\Delta < 0 \end{aligned}$$

which is a contradiction since we constructed a solution better than the optimal solution. \square

4 Euclidean Projection onto the box-constrained L_1 Ball

We modify the problem studied by Duchi et al. [7] to the more generic scenario containing the bounds on the predicted vector. We need to find the projection of a vector $\mathbf{v} \in \mathbb{R}^n$ onto a feasible region defined by

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \text{ s.t. } \{\mathbf{x} \in \Omega : \|\mathbf{x}\|_1 \leq z, \mathbf{a} \preceq \mathbf{x} \preceq \mathbf{b}\} \quad (2)$$

Note that the vector \mathbf{v} is no longer contained within the positive real space. We assume that $\Omega \neq \emptyset$ guarantees feasibility and existence of a solution. The range in which the

x_j 's should lie can now be characterized as (a) $[a_j, b_j] < 0$, (b) $0 < [a_j, b_j]$ and (c) $a_j \leq 0 \leq b_j$, assuming that $a_j < b_j$ for all cases.

Intervals not containing 0: Further analysis of the cost function in Eq.(2), leads to the observation that conditions (a) and (b), are equivalent under a sign flip. This can be obtained by observing the following identities: (a) distance preservation under sign flip $\|\mathbf{v} - \mathbf{x}\|_2^2 = \|(-\mathbf{v}) - (-\mathbf{x})\|_2^2$, (b) 1 norm preservation under sign flip $\|\mathbf{x}\|_1 = \|-\mathbf{x}\|_1$ and (c) range transformation under sign flip $\mathbf{x} \in [\mathbf{a}, \mathbf{b}] < 0 \Leftrightarrow -\mathbf{x} \in [-\mathbf{b}, -\mathbf{a}] > 0$. For such constraints, we can transform the bounds such that all the boundaries are positive. Once this is done a simple change of variables

$$\{\mathbf{v}, \mathbf{x}, \mathbf{b}\} - \mathbf{a} = \hat{\mathbf{v}}, \hat{\mathbf{x}}, \hat{\mathbf{b}}, \quad z - \|\mathbf{a}\|_1 = \hat{z} \quad (3)$$

leads to the formulation in Eq.(2) with the lower bound terms a_i 's equal to 0. The equivalent simpler problem is exactly similar to Eq.(1), since the L_1 constraint and the simplex constraint are same for positive variables. Also note that element wise manipulations can be performed when all the bounds do *not* belong to one particular case, without altering the form of the equations.

Interval containing 0: The objective function in Eq.(2) has expression of the form $\|\mathbf{v} - \mathbf{x}\|_p$ ($p = 2$) and norm constraint is equivalent to inclusion in L_p norm ball ($p = 1$). Given a candidate solution $\mathbf{x} \in \mathbb{R}^n$, let us define a category of moves which can be used to generate a family of candidate solutions. By setting some subset of vector components to zero, we can generate corresponding points in all orthants, resulting in a family of up to 2^n candidate solutions, one in each orthant. We note two properties of orthant projection move which are essential for our exposition. First, orthant projection move preserves L_p norm ball inclusion constraint for all p . Second, under orthant projection moves of \mathbf{x} , $\|\mathbf{v} - \mathbf{x}\|_p$ is minimized when \mathbf{x} and \mathbf{v} are in the same orthant. Together, these two properties ensure that there is a preferred orthant (determined apriori) where optimal solution is guaranteed to lie, as long as none of the other constraints are violated. This observation can be used to generalize Lemma 3 of Duchi et al. [7] to a much wider class of problems. For completion and ease of understanding we state the following lemma:

Lemma 2. [Lemma 3, Duchi et al. [7]] *Let \mathbf{x} be an optimal solution of Eq.(2). Then, $x_i v_i \geq 0$, $\forall i$.*

Hence, x_i has the same sign as v_i . Note that this lemma holds true for the upper bounded problem as well, since replacing any variable with 0 does not violate this constraint, and hence the proof for the lemma can be exactly applied to the more generic case mentioned above.

Orthant projection also reveals a possible failure mode for the above generalization. If there are terms other than those of the form of $\|\mathbf{v} - \mathbf{x}\|_p$ and there are constraints other than those of the form of inclusion in L_p norm ball, then apri-

ori preferred orthant selection may not be possible. Preferred orthant selection allows us to simplify the functional form of L_p norm terms. In particular, for L_1 norm projection problem, once problem has been transformed to guarantee that optimal solution lies in first orthant, L_1 norm constraint becomes equivalent to the simplex constraints, $x_i \geq 0 \forall i$ and $\sum_{i=1}^n x_i \leq z$.

In order to retain this simplification, any generalization involving terms which are not conducive to orthant projection must be explicitly handled. Box constraints are in general not conducive, unless interval contains origin. For intervals such as $a_j < 0 < b_j$, we have the following claim:

$$\textbf{Claim 3. } |x_i| \in \begin{cases} \{0, |a|\} & v_i < 0 \\ \{0, |b|\} & v_i > 0 \end{cases}$$

Based on the above discussions we can write the generic upper bounded L_1 norm projection problem as, given any $\hat{\mathbf{v}} \in \mathbb{R}^n$, take the absolute value of the elements and transform it to $\mathbf{v} \in \mathbb{R}_+^n$, transform $[\mathbf{a}, \mathbf{b}] \Rightarrow [\mathbf{0}, \mathbf{b}]$ based on claim. 3, find \mathbf{x} by solving

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \quad \text{s.t. } \{\mathbf{x} \in \Omega : \|\mathbf{x}\|_1 \leq z, \mathbf{0} \preceq \mathbf{x} \preceq \mathbf{b}\} \quad (4)$$

and return the final projection as $\hat{\mathbf{x}} = \mathbf{x} * \text{sign}(\hat{\mathbf{v}})$.

Identifying the simplex projection as well as L_1 projection as the same problem leads us to study the unified problem mentioned above in Eq.(4).

The Lagrangian for the above optimization problem (Eq.(4)) can be written as

$$\mathcal{L} = \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + \theta \left(\sum_{i=1}^n x_i - z \right) - \zeta \cdot \mathbf{x} - \gamma \cdot (\mathbf{b} - \mathbf{x}) \quad (5)$$

Differentiating with respect to x_i and comparing to zero gives the first order optimality condition,

$$\frac{\partial \mathcal{L}}{\partial x_i} = x_i - v_i + \theta - \zeta_i + \gamma_i = 0 \quad (6)$$

The first complementary slackness KKT condition [2], implies that $x = 0$, when $v_i + \zeta_i = \theta_i$. Since, $\zeta_i > 0$, hence $x_i = 0$ whenever $v_i < \theta$. The second complementary slackness KKT condition implies that $0 < x_i < b_i$, means $\zeta_i = 0, \gamma_i = 0$ and

$$x_i - v_i + \theta = 0 \quad (7)$$

The addition of finite upper bound leads to the next complementary slackness condition, namely, when the value of x_i reaches it maximum value b_i , $\gamma_i > 0$ and

$$b_i - v_i + \theta + \gamma_i = 0 \quad (8)$$

Claim 4. $x_i = b_i$ implies $v_i \geq b_i$.

Note that the converse is not generally true, that is $v_i > b_i$ does not imply that $x_i = b_i$, since it can still be lower than the upper bound.

Corollary 4. $v_i > b_i$ and $\gamma_i > 0$ implies $x_i = b_i$.

One important aspect of the cost $\|\mathbf{x} - \mathbf{v}\|_2^2$ is that the contribution of each x_i to the total cost is dependent on the distance $v_i - x_i$. From this point onwards we will assume that the upper bound term $b_i = \min(v_i, b_i)$, such that $v_i - b_i \geq 0$. Since each x_i is bounded to be less than b_i , hence $v_i - b_i$ can be thought to be the relative weight determining the order in which x_i 's should be changed to meet the norm constraint. This ordering can also be argued from the fact that the magnitude of the gradient with respect to x_i is determined by the quantity $v_i - b_i$.

In Lemma 1, we have shown that even for upper bounded simplex projection problem (after restriction to first orthant), such constraint ordering is possible for constraints $x_i \geq 0$, and is determined by \mathbf{v} . In the next lemma we show that similar constraint ordering is possible for upper bound constraints $x_i \leq b_i$, and is determined by $(\mathbf{v} - \mathbf{b})$ which is one of our key contributions and forms the basis of the proposed efficient algorithm. Based on the above observations we write a modified version of the lemma 2. from Shalev-Shwartz et al. [13].

Lemma 3. Let \mathbf{x} be an optimal solution of Eq.(4). Let i and j be two indices such that $(v_i - b_i) \leq (v_j - b_j)$. If $x_i = b_i$ then $x_j = b_j$ as well.

Proof. From Eq.(8), whenever $x_i = b_i$, then $v_i - b_i = \theta + \gamma_i$ where $\gamma_i > 0$. Hence

$$\begin{aligned} v_i - b_i &> \theta, & \text{since } \gamma_i > 0 \\ \Rightarrow v_j - b_j &\geq v_i - b_i > \theta, & \text{given} \\ \Rightarrow v_j - b_j &> \theta \\ \Rightarrow v_j - b_j &= \theta + \gamma_j, & \text{such that } \gamma_j > 0 \\ \Rightarrow x_j &= b_j & \text{from Corollary 4. } \square \end{aligned}$$

4.1 Worst case strongly linear time algorithm

We now propose an algorithm with strongly linear time worst case complexity which is asymptotically fastest possible. It is based on dependence between θ and z along regularization path. We have already shown that, in optimal solution $x_i = 0$ whenever $v_i < \theta$, and $x_i = b_i$ whenever $v_i - b_i \geq \theta$ and equal to $v_i - \theta$ otherwise. Hence, for any value of θ , variables x_i 's can be divided into three disjoint groups.

$$x_i = \begin{cases} 0 & \text{if } v_i \leq \theta & \text{Fixed at lower limit} \\ b_i & \text{if } v_i - b_i \geq \theta & \text{Fixed at upper limit} \\ v_i - \theta & \text{if } v_i - b_i < \theta < v_i & \text{Constraints inactive} \end{cases} \quad (9)$$

Let us denote the sets of indices of x_i 's in these groups by L , U and C respectively. These sets are functions of θ . Let optimal θ be θ^* and corresponding sets be L^* , U^* and C^* . Relation between z and θ can be expressed as,

$$z = \sum_{i=1}^n x_i = \sum_{i \in U} b_i + \sum_{i \in C} (v_i - \theta) = \sum_{i \in U} b_i + \sum_{i \in C} v_i - |C|\theta \quad (10)$$

where $| \cdot |$ for a set argument, denotes its cardinality. It is evident that z is monotonically decreasing piece-wise linear function of θ , with $2n$ points of discontinuity at v_i and $(v_i - b_i)$ values. The pseudo code of our proposed algorithm is given in Algorithm 1. The algorithm operates upon merged \mathbf{v} and $(\mathbf{v} - \mathbf{b})$ arrays, maintaining source information.

In the first stage, we find the linear segment corresponding to given z . Uncertainty interval $[\theta_L, \theta_R]$ for θ is initialized with $[\min(\mathbf{v} - \mathbf{b}), \max(\mathbf{v})]$ and is subsequently reduced in every iteration by bisection at a pivot selected from the elements of merged \mathbf{v} and $(\mathbf{v} - \mathbf{b})$ arrays lying in current uncertainty interval. For pivot, we use median, found using worst case linear time median finding algorithm [4], in order to ensure that number of iterations remains $O(\log n)$ and that after every iteration, size of uncertainty interval reduces by a constant fraction. Using Eq.(10) to evaluate z at pivot θ is not efficient enough for overall linear time complexity, since summations involve $O(n)$ terms every time, resulting in $O(n \log n)$ complexity. To rectify this inefficiency, apart from $S_{all} = \sum_{i=1}^n v_i$, we maintain two running partial sums across all iterations.

1) S_L = sum of v_i for all elements which are guaranteed to be set to zero in optimal solution i.e. $v_i \leq \theta_L \Rightarrow i \in L^*$.

2) S_R = sum of $(v_i - b_i)$ for all elements which are guaranteed to be set to corresponding upper bounds in optimal solution i.e. $(v_i - b_i) \geq \theta_R \Rightarrow i \in U^*$.

We also maintain cardinality n_L and n_R of these sets. In terms of these partial sums, (10) can be expressed as

$$\begin{aligned} z_{pivot} &= S_{all} - S_L - S_R - (n - n_L - n_R) * \theta_{pivot} \\ &- \sum_{i: \theta_{pivot} \leq v_i - b_i < \theta_R} (v_i - b_i - \theta_{pivot}) \\ &- \sum_{i: \theta_L < v_i \leq \theta_{pivot}} (v_i - \theta_{pivot}) \end{aligned} \quad (11)$$

If $z_{pivot} > z_{target}$, $[\theta_{pivot}, \theta_R]$ becomes new uncertainty interval, and S_L and n_L are updated as

$$\{S_L, n_L\} \leftarrow \{S_L, n_L\} + \sum_{i: \theta_L < v_i \leq \theta_{pivot}} \{v_i, 1\} \quad (12)$$

Otherwise, if $z_{pivot} < z_{target}$, $[\theta_L, \theta_{pivot}]$ becomes new

uncertainty interval, and S_R , and n_R are updated as

$$\{S_R, n_R\} \leftarrow \{S_R, n_R\} + \sum_{i: \theta_{pivot} \leq v_i - b_i < \theta_R} \{(v_i - b_i), 1\} \quad (13)$$

Iterations continue until there are no more points of discontinuity in the uncertainty interval and L^* , U^* and C^* have been found. Now, following modified version of (10) can be used to evaluate θ^* as

$$\theta^* = \frac{\sum_{i \in U^*} b_i + \sum_{i \in C^*} v_i - z_{target}}{|C^*|} \quad (14)$$

Algorithm 1 Algorithm for worst case strongly linear time projection onto the simplex with finite upper bound.

```

REQUIRE  $\mathbf{v} \in \mathbb{R}^n$ ,  $\mathbf{b} \in \mathbb{R}^n$ ,  $0 < z_{target} < \sum_{i=1}^n b_i$ 
 $\mathbf{v}\mathbf{b} \leftarrow \text{merge}(\mathbf{v}, (\mathbf{v} - \mathbf{b}))$  //maintain source info
 $(idx\_theta_L, idx\_theta_R) \leftarrow (0, 2n - 1)$ 
 $S_{all} \leftarrow \text{sum}(\mathbf{v})$ 
 $(n_L, S_L) \leftarrow 0$ 
 $(n_R, S_R) \leftarrow 0$ 
while  $idx\_theta_R > idx\_theta_L + 1$  do
   $\theta_{pivot} \leftarrow \text{pivot\_select}(\mathbf{v}\mathbf{b}, idx\_theta_L, idx\_theta_R)$ 
   $\text{partition}(\mathbf{v}\mathbf{b}, idx\_theta_L, idx\_theta_R, \theta_{pivot})$ 
   $idx\_theta_{pivot} \leftarrow \text{index}(\theta_{pivot})$ 
  Evaluate  $z_{pivot}$  using (11)
  if  $z_{pivot} > z_{target}$  then
     $idx\_theta_L \leftarrow idx\_theta_{pivot}$ 
    Update  $(S_L, n_L)$  using (12)
  else
     $idx\_theta_R \leftarrow idx\_theta_{pivot}$ 
    Update  $(S_R, n_R)$  using (13)
  end if
end while
Evaluate  $\theta^*$  using (14)
RETURN  $\mathbf{x}$  corresponding to  $\theta^*$  (Eq.9).
```

5 Experiments

The first set of experiments are performed for synthetic data. We generate labeled data belonging to 2 classes such that the probability of the label being 1/0 is distributed according to logistic likelihood $p(y_i = 1 | \mathbf{x}_i, \mathbf{w}) = \sigma(\mathbf{w} \cdot \mathbf{x}_i)$, where $\sigma(a) = 1/(1 + \exp(-a))$. Additionally, we disturb 10% of the data labels by introducing false labels based on random draws. The ground truth parameter vector \mathbf{w} is generated from a generalized Gaussian distribution, with rejection, such that the individual elements of the vector are bounded within ± 0.5 . Moreover half the entities of \mathbf{w} are made zeros to generate a sparse vector. The inference problem is thus an estimation problem with known upper bounds. We minimize the average logistic log loss, and project the gradient vector to the convex space. The norm

constraint is determined as a fixed fraction of the dimension of the vector. The estimation error against the iterations, where the error is denoted as $f(\mathbf{w}) - f(\mathbf{w}^*)$, for L1 and our method called UB.L1 is shown in Fig. 3(left). Note that the L1 method estimates are outside the bounds which manifests itself as slower rate of convergence as evident from the plots. At convergence UB.L1 estimate seems much closer to the ground truth parameter vector than the L1 estimate.

Next we explore the run time performance of our algorithm, against a standard quadratic programming (QP) implementation in MATLAB. Fig. 4 shows the results for such an experiment. For projecting dense vectors with 1M non-zero elements our method takes around 0.22 seconds.

Food distribution The next experiment is drawn from a real world scenario and motivates the upper bound constraints. We start by noting that the problem of food distribution can be easily applied to our case. Suppose the *production* of one food item (e.g. chicken) in 40 states in the US is provided as the initial vector \mathbf{v}^1 . Assume that $r\%$ of the total production is put up for sales.

We would like to find the *sales* vector \mathbf{x} for the 40 states. The upper bounds can be obtained from the consumption patterns in the previous years. We take production in 2007 as the new vector \mathbf{v} , and the distribution in 2004 as the upper bound \mathbf{b} . To remove scale differences the upper bound is normalized such that $\|\mathbf{b}\|_2 = \|\mathbf{v}\|_2$. The norm constraint $z = \|\mathbf{v}\|_2 * r/100$. The results for such an experiment are shown in Fig. 5. As the value of z decreases L1 forces more and more mass into the dominant elements. Our method still tries to satisfy the upper bound constraints, which spreads the distribution at the cost of sparsity. As the supply decreases, L1 tries to bias the distributions among the states based on the relative weights of the production itself. Our method, on the other hand, applies the demand based upper bounds, and biases the distribution in favor of the states with *maximum disparity* between production and supply.

6 Conclusion

In this paper we extend the idea of L_1 constrained gradient projection under the presence of upper bound constraints. We explore simplex projection with upper bounds and bring out the similarities with L_1 projection. We derive criteria for a-priori determination of sequence in which various constraints become active and use such orderings to propose an efficient algorithm. The key insight obtained from our experiments was that L_1 tries to increase the dominant elements while putting zeroes for all the others. Bound constrained L_1 , weighs the elements based on their distance from the corresponding bound. This case leads

¹<http://www.agcensus.usda.gov/Publications/2007>




Figure 3: Top left: plot for $\text{abs}(f(\mathbf{w}) - f(\mathbf{w}^*))$ vs. iteration (training). Top right: similar plot for testing. Bottom: estimated parameter vector \mathbf{w} . For all panels, blue: L1, red: our method UB_L1 and green: ground truth.




Figure 4: Left: comparison of our method against Matlab QP. Blue: run time (in seconds) for QP implementation of MATLAB. Red: run time for our linear time method. The horizontal axis runs over the dimension of the input vector v . Right: zoomed in red curve.




Figure 5: Red: Actual sales of Chicken in 40 states in 2007. Blue: our method with upper bounds. Green: L_1 only. Note the small region in the circle which has been enlarged. L_1 completely misses this region whereas our method still provides some value to it.

to better predictions, specifically in cases which should be weighed based on the disparity between the demand and supply. The elements with higher disparity get higher weight in the predicted distribution vector.

References

- [1] D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.
- [2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [3] E. J. Candes. Compressive sampling. *Proc. Intl. Congress of Mathematicians*, 2006.
- [4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press and McGraw-Hill, second edition, 2001.
- [5] Koby Crammer and Yoram Singer. On the learnability and design of output codes for multiclass problems. In *Proc. 13th Annual Conf. Computational Learning Theory*, pages 35–46, 2000.
- [6] David L. Donoho. For most large underdetermined systems of linear equations the minimal l_1 -norm solution is also the sparsest solution. *Comm. Pure Appl. Math*, 59:797–829, 2004.
- [7] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l_1 -ball for learning in high dimensions. In *ICML '08: Proc. 25th Intl. Conf. Machine learning*, 2008.
- [8] E. Gafni and D. P. Bertsekas. Two-metric projection methods for constrained optimization. In *SIAM J. Control and Optimization*, pages 936–964, 1984.
- [9] Kwangmoo Koh, Seung-Jean Kim, and Stephen Boyd. An interior-point method for large-scale l_1 -regularized logistic regression. *J. Mach. Learn. Res.*, 8, 2007.
- [10] Nimrod Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31(1):114–127, 1984.
- [11] Nimrod Megiddo and Arie Tamir. Linear time algorithms for some separable quadratic programming problems. *Operations Research Letters*, 13:203–211, 1993.
- [12] A. Y. Ng. Feature selection, l_1 vs. l_2 regularization, and rotational invariance. In *ICML '04: Proc. Twenty-first Intl. Conf. on Machine learning*, 2004.
- [13] S. Shalev-Shwartz and Y. Singer. Efficient learning of label ranking by soft projections onto polyhedra. In *J. Mach. Learn. Res.*, volume 7, pages 1567–1599, July 2006.
- [14] Charles M. Stein. Estimation of the mean of a multivariate normal distribution. In *The Annals of Statistics*, volume 9, pages 1135–1151, Nov 1981.
- [15] Robert Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal Statistical Society, Series B*, 58:267–288, 1994.
- [16] Yaakov Tsaig and David L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52:1289–1306, 2006.
- [17] Hui Zou, Trevor Hastie, and Robert Tibshirani. On the “degrees of freedom” of the lasso. Dec 2007.